

WHAT IS CLAIMED IS:

1. An apparatus comprising:

5 a computing node configured to generate a request to open a file, wherein said
 computing node is configured not to cache data from said file; and

 a metadata server coupled to receive said request, wherein said metadata server is
 configured to provide at least a first file identifier corresponding to said
10 file responsive to said request, said first file identifier identifying at least a
 portion of said file within a storage storing said file;

 wherein said computing node is coupled to receive said first file identifier for
15 directly accessing said storage.

2. The apparatus as recited in claim 1 wherein said metadata server is further configured
to provide a device identifier indicative of said storage responsive to said request. }

3. The apparatus as recited in claim 2 wherein said computing node is coupled to receive
20 said device identifier, and wherein said computing node is configured to route accesses to
 said storage using said device identifier.

4. The apparatus as recited in claim 1 wherein said metadata server is configured to
assign an access key to said request.

25 5. The apparatus as recited in claim 4 wherein said computing node is coupled to receive
 said access key, and wherein said computing node is configured to transmit said access
 key with an access command to said storage.

6. The apparatus as recited in claim 5 further comprising said storage, wherein said storage is coupled to receive said access key from said metadata server, and wherein said storage is configured to validate said access key from said computing node.

5 7. The apparatus as recited in claim 1 further comprising said storage, wherein said computing node is configured to transmit an access command to said storage to close said file, and wherein said storage is configured to transmit an indication to said metadata server including said first file identifier to inform said metadata server of said file being closed.

10

8. The apparatus as recited in claim 1 wherein said request includes a file name of said file.

15

9. The apparatus as recited in claim 8 wherein said metadata server includes a directory, and wherein said directory maps said file name to said first file identifier.

20

10. The apparatus as recited in claim 1 further comprising said storage, wherein said computing node is configured to transmit an access command to said storage to access one or more bytes, and wherein a number of said one or more bytes affected by said access command is not equal to a block size of said storage.

25

11. The apparatus as recited in claim 1 further comprising said storage and an interconnect, wherein said computing node, said metadata server, and said storage are each coupled directly to said interconnect.

12. The apparatus as recited in claim 11 wherein said computing node is configured to transmit said request to said metadata server through said interconnect, and wherein said computing node is configured to receive said first file identifier from said metadata server through said interconnect.

13. The apparatus as recited in claim 11 wherein said computing node is configured to directly access said storage through said interconnect.

5 14. The apparatus as recited in claim 1 further comprising said storage, wherein said computing node is configured to transmit an access command to said metadata server to close said file, and wherein said metadata server is configured to transmit an indication to said storage including said first file identifier to inform said storage of said file being closed.

10

15. A method comprising:

generating a request to open a file from a computing node;

15

providing at least a first file identifier corresponding to said file from a metadata server responsive to said generating said request, said first file identifier corresponding to said file and identifying said file within a storage storing at least a portion of said file; and

20

directly accessing said storage from said computing node responsive to said first file identifier, wherein said computing node is configured not to cache data from said file.

16. The method as recited in claim 15 further comprising providing a device identifier
25 identifying said storage from said metadata server responsive to said generating said request.

17. The method as recited in claim 16 further comprising routing an access from said computing node directly to said storage using said device identifier.

18. The method as recited in claim 15 further comprising assigning an access key to said computing node responsive to said generating said request.

5 19. The method as recited in claim 18 further comprising transmitting said access key from said computing node to said storage with an access command for said storage.

20. The method as recited in claim 19 further comprising:

10 receiving said access key from said metadata server into said storage;

receiving said access key from said computing node into said storage; and

15 validating said access key from said computing node using said access key from said metadata server.

21. The method as recited in claim 15 wherein said directly accessing comprises generating an access command for one or more bytes and transmitting said access command to said storage, wherein a number of said one or more bytes does not equal a
20 block size of said storage.

22. An apparatus comprising:

25 a computing node; and

a low latency interconnect coupled to the computing node and to a storage,
wherein the computing node is configured to directly access the storage
via the low latency interconnect using a first file identifier identifying a
file on said storage.

23. The apparatus as recited in claim 22 wherein the low latency interconnect has a typical latency less than 100 microseconds.

5 24. The apparatus as recited in claim 22 further comprising a metadata server coupled to the low latency interconnect, wherein said computing node is configured to generate a request to open a file, and wherein said metadata server is coupled to receive said request, wherein said metadata server is configured to provide at least said first file identifier corresponding to said file responsive to said request, said first file identifier identifying at
10 least a portion of said file within said storage, and wherein said computing node is coupled to receive said first file identifier for directly accessing said storage.

25. A method comprising:

15 directly accessing a storage from a computing node coupled to the storage via a low latency interconnect, the directly accessing responsive to a first file identifier.

20 26. The method as recited in claim 25 wherein the low latency interconnect has a typical latency less than 100 microseconds.

27. The method as recited in claim 25 further comprising:

generating a request to open said file from said computing node;
25 providing at least said first file identifier corresponding to said file from a metadata server responsive to said generating said request, said first file identifier corresponding to said file and identifying said file within said storage.

28. An apparatus comprising:

5 a first computing node configured to directly access a storage with a first access command;

a second computing node configured to directly access said storage with a second access command concurrent with said first computing node; and

10 said storage coupled to said first computing node and said second computing node.

29. The apparatus as recited in claim 28 further comprising a metadata server, wherein said first computing node and said second computing node are configured to directly
15 access said storage with said first access command and said second access command independent of said metadata server.

30. A method comprising:

20 directly accessing a storage with a first access command from a first computing node;

directly accessing said storage with a second access command from a second computing node concurrent with said directly accessing said storage with
25 said first access command.

31. The method as recited in claim 30 wherein said directly accessing said storage with said first command and said directly accessing said storage with said second access command are independent of a metadata server.

30